Journal of Statistical Mechanics: Theory and Experiment
An IOP and SISSA journal

PAPER: STATPHYS 28

Evolutionary accessibility of random and structured fitness landscapes

Joachim Krug* **and Daniel Oros**

Institute for Biological Physics, University of Cologne, Zülpicher Strasse 77, 50937 Köln, Germany
E-mail: jkrug@uni-koeln.de

Received 29 November 2023
Accepted for publication 28 February 2024
Published 3 April 2024

Online at stacks.iop.org/JSTAT/2024/034003
https://doi.org/10.1088/1742-5468/ad3197

Abstract. Biological evolution can be conceptualized as a search process in the space of gene sequences guided by the fitness landscape, a mapping that assigns a measure of reproductive value to each genotype. Here, we discuss probabilistic models of fitness landscapes with a focus on their evolutionary accessibility, where a path in a fitness landscape is said to be accessible if the fitness values encountered along the path increase monotonically. For uncorrelated (random) landscapes with independent and identically distributed fitness values, the probability of existence of accessible paths between genotypes at a distance linear in the sequence length L becomes nonzero at a nontrivial threshold value of the fitness difference between the initial and final genotypes, which can be explicitly computed for large classes of genotype graphs. The behaviour of uncorrelated random landscapes is contrasted with landscape models that display additional, biologically motivated structural features. In particular, landscapes defined by a tradeoff between adaptation to environmental extremes have been found to display a combinatorially large number of accessible paths to all local fitness maxima. We show that this property is characteristic of a broad class of models that satisfy a certain global constraint, and provide further examples from this class.

Keywords: evolution models, evolutionary processes, energy landscapes, genetics

*Author to whom any correspondence should be addressed.

J. Stat. Mech. (2024) 034003

Contents

1. Introduction

Klettere, steige, steige. Aber es gibt keine Spitze[1]
Kurt Tucholsky

The description of biological evolution in terms of a fitness landscape was introduced by Sewall Wright almost a century ago [1]. The underlying conceptual step is as simple as it is bold: Given that evolution is driven by heritable variation within a population, where the genes of individuals with better reproductive capability (*fitness*) spread by natural selection, it is tempting to consider a theoretical framework that maps the

[1] Climb, ascend, ascend. But there is no summit.

J. Stat. Mech. (2024) 034003

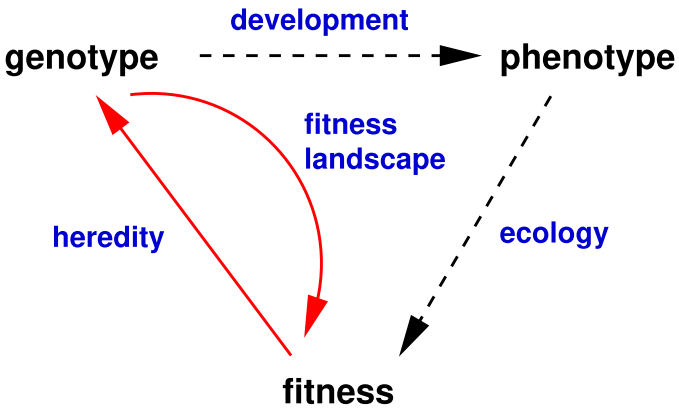

Figure 1. The fitness landscape concept in evolutionary theory. Darwinian evolution is driven by a feedback cycle consisting of three major steps. Through the process of *development*, the physical features of the organism (its *phenotype*) emerge from the genotype. *Ecological interactions* of the phenotype with the abiotic and biotic environment determine the organism's reproductive success (its *fitness*). Finally, by the phenomenon of *heredity*, the reproductive success determines the relative abundance of genotypes in the next generation. The fitness landscape constitutes a conceptual shortcut where the complexities of development and ecology are bypassed. The figure was inspired by a lecture by Amitabh Joshi at the Jawaharlal Nehru Centre for Advanced Scientific Research in Bangalore in 2015.

genetic makeup of an organism (its *genotype*) directly to its fitness. Needless to say, this is an enormous oversimplification, by which much of the structure of the biological world is swept under the rug (figure 1). Nevertheless, the fitness landscape concept has proven to be of great heuristic value, and it continues to play a central role in current evolutionary theory [2–8].

Already in his seminal article, Wright expressed the concern that the *ruggedness* of the fitness landscape, that is, the existence of multiple fitness peaks, might compromise the ability of natural selection to access optimal evolutionary solutions:

> *In a rugged field of this character, selection will easily carry the species to the nearest peak, but there will be innumerable other peaks that will be higher but which are separated by 'valleys'. The problem of evolution as I see it is that of a mechanism by which the species may continually find its way from lower to higher peaks in such a field.* [1]

Wright's expectation that real fitness landscapes possess *'innumerable other peaks'* was not shared by everyone, however. In a letter written in 1931, Ronald Fisher speculated that true fitness peaks should in fact be very rare, because most stationary points of a function on a high-dimensional space are saddles rather than maxima or minima [9]. Similar arguments can be found in the recent literature [10], and the question to what extent our view of the evolutionary process is misled by intuition gained from low-dimensional landscapes is still under debate [11].

J. Stat. Mech. (2024) 034003

In this contribution we address the evolutionary accessibility of high-dimensional rugged fitness landscapes within a somewhat restricted, but mathematically clearly defined setting. Formalizing the fitness landscape as a real-valued function on the space of genetic sequences, we ask for the probability that two of these sequences are connected by a mutational path along which the fitness function increases monotonically [12–15]. This notion of accessibility was first formulated by Daniel Weinreich and collaborators in their pioneering studies of small combinatorially complete fitness landscapes describing the evolution of antimicrobial drug resistance [16–18].

Meanwhile, the scale of the empirical fitness landscapes that can be explored using state-of-the-art high-throughput methods has increased enormously [19–27], but the factors determining how accessible these landscapes are remain largely unknown. For example, in an analysis of the affinity landscape spanned by the 15 point mutations separating the spike protein of the ancestral Wuhan strain of the SARS-CoV2 virus from the Omicron BA.1 variant, it was found that none of the $15! \approx 1.3 \times 10^{12}$ direct mutational pathways[2] are monotonically increasing [25]. On the other hand, a study of combinations of nucleotides at nine positions of the gene coding for the dihydrofolate reductase (DHFR) enzyme in *Escherichia coli* subjected to the antibiotic trimethoprim revealed a fitness landscape that combines high ruggedness with high accessibility [26]; we will return to this example below in section 5.

The purpose of this article is to provide a concise and coherent account of our current understanding of evolutionary accessibility in probabilistic models of fitness landscapes. We distinguish between *random* (uncorrelated) landscapes, where fitness values assigned to genotypes are independent and identically distributed (i.i.d.) random variables, and *structured* landscapes that possess some kind of (more or less biologically motivated) fitness correlation. The next section introduces the mathematical setting, and explains in particular the description of sequence space as the Cartesian power of a mutation graph [28]. Known results for random and structured landscapes are reviewed in sections 3 and 4, respectively. An important lesson from the work on structured landscapes is that ruggedness does not generally correlate with accessibility in a simple way. Section 5 describes a recently discovered class of fitness landscape models that combine high accessibility with high ruggedness. Most of the material in this section has not been reported previously. The article concludes with a summary and an outlook in section 6.

2. Sequence space and accessible paths

A genotype is a sequence $\sigma = (\sigma_1, \ldots, \sigma_L)$ of length L with entries σ_i drawn from an alphabet of size a, which we represent by the numbers $0, \ldots, a-1$. In the following, the entries in the sequence will be referred to as *alleles*, and the positions $i = 1, \ldots, L$ in the sequence as *loci*. The set $\{0, \ldots, a-1\}^L$ equipped with the natural distance measure

$$d(\sigma, \tau) = \sum_{i=1}^{L} 1 - \delta_{\sigma_i, \tau_i} \tag{1}$$

[2] See section 2 for the definition of direct paths.

J. Stat. Mech. (2024) 034003

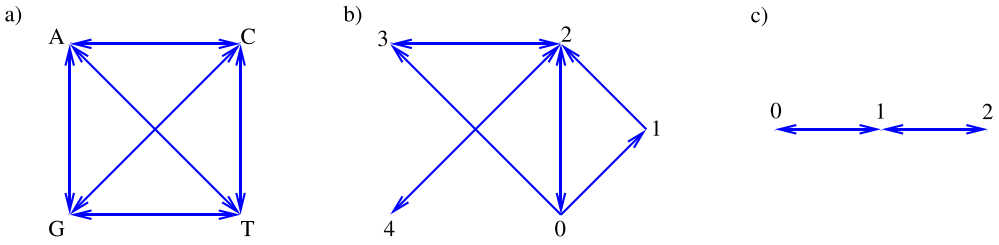

Figure 2. Examples of allele graphs. (a) The allele graph of the nucleotide alphabet of DNA sequences is the complete graph over four alleles. (b) An incomplete graph over five alleles. This graph is the smallest known example that admits an accessibility problem of irregular type. In this case the paths are assumed to move from $\alpha = (0 \ldots 0)$ to $\omega = (4 \ldots 4)$. See section 3.2 for details. (c) The linear path graph over three alleles. Such a graph appears naturally if the different alleles represent the copy number of a gene, which changes by deletion or duplication [30].

defines the *Hamming graph* $\mathbb{H}_a^L$ [29], where genotypes at a mutational distance $d = 1$ are connected by a link. The binary Hamming graph $\mathbb{H}_2^L$ is the L-dimensional discrete hypercube.

The description of sequence spaces by Hamming graphs assumes that any allele can mutate to any other allele, which is not always the case. For example, the genetic code restricts the possible single-nucleotide mutations that convert one amino acid into another. To account for this, it is useful to introduce the *allele graph* $\mathcal{A}$ over the allele set $\{0, \ldots, a-1\}$, which encodes the allowed mutational transitions [28]. The adjacency matrix $\mathbf{A} = \{A_{\mu\nu}\}_{\mu,\nu=0,\ldots,a-1}$ of $\mathcal{A}$ has an entry $A_{\mu\nu} = 1$ iff the mutation $\mu \to \nu$ is possible, and $A_{\mu\nu} = 0$ otherwise. Examples of allele graphs are shown in figure 2. The sequence space is then the L-fold Cartesian product of $\mathcal{A}$, which we denote by $\mathcal{A}^L$. The Hamming graph $\mathbb{H}_a^L$ is the L-fold product of the complete graph over a alleles.

A fitness landscape is a real-valued function

$$g : \mathcal{A}^L \to \mathbb{R}. \tag{2}$$

By decorating the links by arrows pointing in the direction of increasing fitness, the fitness landscape induces an *orientation* on the graph, which by construction is acyclic (figure 3). The resulting oriented graph is called a *fitness graph* [31, 32].

In order to make the fitness graph construction unambiguous, we will assume throughout that g is non-degenerate, such that no two genotypes have exactly the same fitness. While this condition is naturally satisfied in probabilistic models where fitness values are drawn from a continuous distribution, it may seem to ignore the important role of neutral mutations [7, 11]. In this context it should be noted, however, that, empirically, not even synonymous mutations (i.e. mutations that change the nucleotide in the DNA sequence but not the corresponding amino acid) are strictly neutral [33, 34].

A *mutational pathway* on $\mathcal{A}^L$ is a sequence of genotypes

$$\alpha = \sigma^{(0)} \to \sigma^{(1)} \to \cdots \to \sigma^{(\ell)} = \omega \tag{3}$$

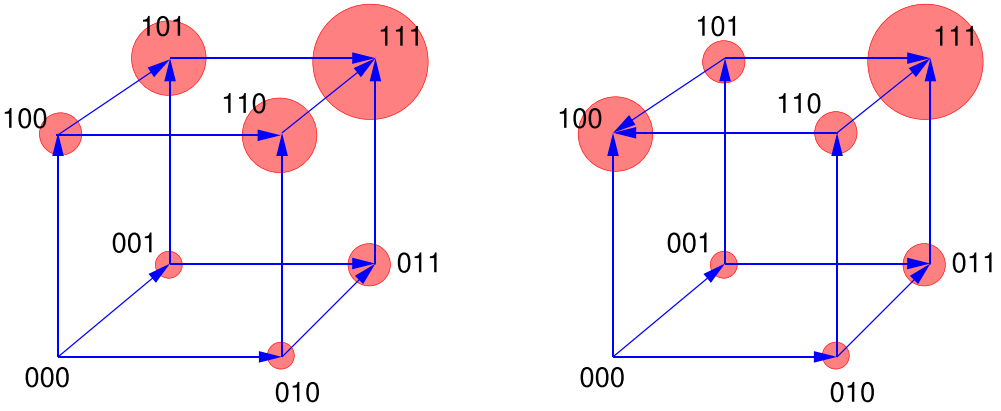

Figure 3. Fitness graphs on the 3-cube. The figure shows two examples of fitness landscapes on the binary 3-cube. Fitness values are represented by the size of the circles on the nodes, and the arrows on the links point in the direction of increasing fitness. The landscape in the left panel has a unique maximum at $\sigma = (1,1,1)$, and there are six direct accessible paths of minimal length 3 from $(0,0,0)$ to $(1,1,1)$. In the right panel there is a local maximum at $(1,0,0)$ and a global maximum at $(1,1,1)$. As a consequence, there are now four direct accessible paths $(0,0,0) \to (1,1,1)$, and one direct and two indirect accessible paths $(0,0,0) \to (1,0,0)$. Each of the two indirect accessible paths contains a backstep (mutational reversion) $1 \to 0$.

where $d(\sigma^{(i)}, \sigma^{(i+1)}) = 1$ for $i = 0, \ldots, \ell - 1$, and all mutational steps are allowed by the allele graph. A path is called *accessible* if the fitness increases in each step,

$$g\left(\sigma^{(i+1)}\right) > g\left(\sigma^{(i)}\right) \tag{4}$$

for all i. In population genetic terminology, all mutational steps are *beneficial*, whereas *deleterious* mutations (which decrease fitness) do not occur. In other words, accessible paths respect the orientation of the fitness graph. Because the orientation is acyclic, accessible paths are self-avoiding. A path is called *direct* if

$$d\left(\alpha, \sigma^{(i)}\right) = i, \tag{5}$$

which implies that each step reduces the distance to the target genotype ω by one. Paths which violate (5) are *indirect*.

3. Random fitness landscapes: accessibility percolation

The most basic probabilistic fitness landscape is the House-of-Cards (HoC) model, where fitness values assigned to genotypes are independent, identically and continuously distributed random variables [35, 36]. The HoC model is the conceptual analogue of the Random Energy Model of spin glasses [37]. Since the fitness graph is fully determined by the rank ordering of the fitness values [38], the statistics of accessible paths are independent of the choice of the fitness distribution, provided it is continuous. For

J. Stat. Mech. (2024) 034003

notational convenience, we will nevertheless assume that the fitness values are uniformly distributed on the unit interval $[0,1]$.

For a pair of genotypes α,ω on the HoC landscape, the number of accessible paths from α to ω is a non-negative integer valued random variable $X_{\alpha,\omega}$, and we say that ω is accessible from α if $X_{\alpha,\omega}\geqslant 1$ [15]. The probability for this event obviously depends on the fitness difference

$$\beta = g(\omega) - g(\alpha) \tag{6}$$

and on the distance $d(\alpha,\omega)$. In this section we are concerned with the behaviour of the probability $\mathbb{P}[X_{\alpha,\omega}\geqslant 1]$ as a function of β for paths with $d(\alpha,\omega)\sim L$ when $L\to\infty$. In many cases (to be specified below), this probability displays a sharp transition from 0 to a nonzero value at a threshold β_c, a phenomenon that has been referred to as *accessibility percolation* [39]. Since comprehensive recent accounts of this topic are available [28, 40, 41], we restrict ourselves to a summary of the main ideas and results.

3.1. Direct paths

Consider a path (direct or indirect) of length ℓ that connects two genotypes α and ω with a fitness difference $\beta>0$. The path is accessible if (i) the $\ell-1$ fitness values of the intermediate genotypes $\sigma^{(1)},\ldots,\sigma^{(\ell-1)}$ are in the interval $(g(\alpha),g(\omega))$, which is true with probability $\beta^{\ell-1}$ and (ii) these fitness values are increasingly ordered, which is true with probability $\frac{1}{(\ell-1)!}$. The joint probability of the two events is

$$P_{\beta,\ell} = \frac{\beta^{\ell-1}}{(\ell-1)!}. \tag{7}$$

To highlight the relation to percolation phenomena, we provide an alternative interpretation of this expression [28, 40]. Suppose that we fix the fitness values of the path endpoints at their extremal values, $g(\alpha)=0$ and $g(\omega)=1$, and remove all other genotypes independently with probability $1-\beta$. Then, the probability for the existence of an accessible path *on extant genotypes* is again given by (7), where the factor $\beta^{\ell-1}$ is now the probability that none of the intermediate genotypes have been removed.

For a direct path the sequences $\alpha=(\alpha_1,\ldots,\alpha_L)$ and $\omega=(\omega_1,\ldots,\omega_L)$ differ at ℓ positions, and in each step, the allele α_i at some position i is replaced by the allele ω_i. This implies that the path is confined to a binary hypercube of dimension ℓ [42], and the properties of the allele graph $\mathcal{A}$ are irrelevant (apart from the fact that all mutations on the path have to be allowed). Because there are $\ell!$ paths in total, the expected number of direct accessible paths is

$$\mathbb{E}^{\mathrm{dir}}[X_{\alpha,\omega}] = \ell! P_{\beta,\ell} = \ell\beta^{\ell-1}, \tag{8}$$

which vanishes for $\ell\to\infty$ whenever $\beta<1$. By Markov's inequality

$$\mathbb{P}[X_{\alpha,\omega}\geqslant 1] = \sum_{k=1}^{\infty}\mathbb{P}[X_{\alpha,\omega}=k] \leqslant \sum_{k=1}^{\infty} k\mathbb{P}[X_{\alpha,\omega}=k] = \mathbb{E}[X_{\alpha,\omega}] \tag{9}$$

J. Stat. Mech. (2024) 034003

we conclude that long direct accessible paths are possible only at the maximal fitness difference $\beta = 1$. A more refined analysis [43] shows that the transition to high accessibility occurs at an ℓ-dependent threshold

$$\beta_c(\ell) = 1 - \frac{\ln \ell}{\ell} \tag{10}$$

in the sense that, for a sequence $\{\beta_\ell\}$ of ℓ-dependent fitness differences,

$$\lim_{\ell\to\infty} \mathbb{P}[X_{\alpha,\omega} \geqslant 1 | \beta = \beta_\ell] = \begin{cases} 0 & \beta_\ell < \beta_c(\ell) \\ 1 & \beta_\ell > \beta_c(\ell). \end{cases} \tag{11}$$

We note for later reference that the threshold (10) can be read off from the expected number of paths (8) by demanding that

$$\lim_{\ell\to\infty} \mathbb{E}^{\mathrm{dir}}[X_{\alpha,\omega} | \beta_\ell = \beta_c(\ell)] = 1. \tag{12}$$

The full distribution of the number of direct accessible paths was found in [44].

Instead of constraining the fitness difference between the endpoints of the path, one may also condition the final point to be of high fitness $g(\omega) = 1$, and leave the fitness of the starting point α undetermined [14]. Averaging (8) over the uniformly distributed fitness of the starting point, one obtains

$$\overline{\mathbb{E}^{\mathrm{dir}}[X_{\alpha,\omega}]} = \int_0^1 d\beta\, \ell \beta^{\ell-1} = 1 \tag{13}$$

independent of the path length ℓ. A combinatorial derivation of this result was given in [15].

3.2. Indirect paths

An indirect path is longer than the distance between its endpoints, $\ell > d(\alpha,\omega)$, which reduces the probability (7) for such a path to be accessible. This is, however, offset by an enormous increase in the number of paths [17, 19, 21, 42]. The total number of self-avoiding paths on the L-cube grows double-exponentially in L and is known explicitly only up to $L = 5$ [45]. As will be shown in this section, the enhancement of accessibility by indirect paths leads to the emergence of a nontrivial accessibility threshold $\beta_c \in (0,1)$.

For general allele graphs $\mathcal{A}$ and $a > 2$, the mutual accessibility of two genotypes along general (indirect) paths depends not only on their distance, but also on their allelic composition. Asymptotically, for a large path length $\ell \to \infty$, the relevant properties of the end points α,ω are encoded in the *divergence matrix* $\mathbf{P} = \{p_{\mu\nu}\}_{\mu,\nu=0,\ldots,a-1}$, where $p_{\mu\nu}$ is the fraction of sites $i \in \{1,\ldots,L\}$ at which $\alpha_i = \mu$ and $\omega_i = \nu$. The rescaled Hamming distance between the two genotypes is then given by

$$\delta \equiv \lim_{L\to\infty} \frac{d(\alpha_L,\omega_L)}{L} = 1 - \sum_{\mu=0}^{a-1} p_{\mu\mu}, \tag{14}$$

J. Stat. Mech. (2024) 034003

where $\{(\alpha_L, \omega_L) \in \mathcal{A}^L \times \mathcal{A}^L\}_{L\in\mathbb{N}}$ is a sequence of endpoints whose asymptotic divergence matrix is given by $\mathbf{P}$.

Similar to the direct path problem in section 3.1, we approach the accessibility question through an analysis of the expected number of accessible paths. In principle, this amounts to enumerating all self-avoiding paths of a given length ℓ, multiplying this number with the weight (7), and summing over ℓ [45]. The difficulty of enumerating all self-avoiding paths can, however, be eliminated by introducing the notion of *quasi-accessibility* [28]. The accessibility problem is modified by allowing for multiple visits to the same genotype (thus removing the constraint of self-avoidance), but assigning a new i.i.d. fitness value to it at each visit. In the resulting *extended HoC model,* each genotype is thus equipped with a countable infinite sequence of fitness values that are evaluated in consecutive visits. A (generally self-intersecting) path on the genotype space $\mathcal{A}^L$ is then called quasi-accessible if the fitness values encountered in the extended HoC landscape are monotonically increasing. Denoting the number of quasi-accessible paths between genotypes α and ω by $\tilde{X}_{\alpha,\omega}$, it is proved in [28] that

$$\mathbb{P}\left[\tilde{X}_{\alpha,\omega} \geqslant 1\right] = \mathbb{P}\left[X_{\alpha,\omega} \geqslant 1\right], \tag{15}$$

and therefore the analysis of quasi-accessible paths is sufficient to determine the accessibility in the original model.

A relatively straight-forward calculation shows that the exponential growth rate of the expected number of quasi-accessible paths is given by [28]

$$\Gamma(\beta) \equiv \lim_{L\to\infty} \frac{\ln \mathbb{E}\left[\tilde{X}_{\alpha_L,\omega_L}\right]}{L} = \sum_{\mu,\nu=0}^{a-1} p_{\mu\nu} \ln\left[\left(e^{\beta\mathbf{A}}\right)_{\mu,\nu}\right], \tag{16}$$

where $\left(e^{\beta\mathbf{A}}\right)_{\mu,\nu}$ is an element of the matrix exponential $e^{\beta\mathbf{A}}$. Under certain natural conditions on the matrices $\mathbf{A}$ and $\mathbf{P}$, which are specified in [28], $\Gamma(\beta)$ is an increasing function with a unique zero, which we denote by β^*. For $\beta < \beta^*$, the function $\Gamma(\beta)$ is negative. Thus, the expected number of quasi-accessible paths tends to zero asymptotically, which implies by Markov's inequality (9) that no quasi-accessible paths exist. By (15), this statement extends to the original accessibility problem, and we conclude that β^* is a lower bound on the accessibility percolation threshold β_c. In particular, if $\beta^* > 1$, accessible paths between the two endpoint genotypes cannot exist for any fitness difference $\beta \in (0, 1]$. An example for which this can be shown to be the case is the path graph over $a \geqslant 3$ alleles (see figure 2(c)).

Determining whether the bound obtained from Markov's inequality is tight requires a careful analysis of higher moments of $\tilde{X}_{\alpha\omega}$ [28]. The result of this analysis is a negativity condition on a certain function constructed from $\mathbf{A}$ and $\mathbf{P}$, which was previously introduced in the context of first passage percolation [46]. Accessibility problems that satisfy this condition are said to be of *regular type* [28]. An example of an allele graph for which the negativity condition is violated, and for which the true accessibility threshold

J. Stat. Mech. (2024) 034003

β_c provably exceeds the lower bound $\beta^* < 1$, is shown in figure 2(b). For problems of regular type, it is proved in [28] that

$$\lim_{L\to\infty} \mathbb{P}\left[X_{\alpha_L,\omega_L} \geqslant 1 | \beta > \beta^*\right] \geqslant C \tag{17}$$

for some positive constant C. Based on related results in [46] and the numerical simulations reported in [42], it is conjectured in [28] that the bound can actually be improved to its maximal value $C = 1$. Assuming this to be the case, we can summarize the percolation behaviour for accessibility problems of regular type in the form

$$\lim_{L\to\infty} \mathbb{P}[X_{\alpha_L,\omega_L} \geqslant 1 | \beta] = \begin{cases} 0 & \beta < \beta_c \\ 1 & \beta > \beta_c, \end{cases} \tag{18}$$

where β_c is the solution of $\Gamma(\beta_c) = 0$. Moreover, the typical length ℓ^* of accessible paths near the threshold is given to leading order in L by

$$\lim_{L\to\infty} \frac{\ell^*(\alpha_L,\omega_L)}{L} = \beta_c \Gamma'(\beta_c) \tag{19}$$

where $\Gamma' = \frac{d\Gamma}{d\beta}$ [28].

3.3. Accessibility of the Hamming graphs

We recall that the Hamming graph $\mathbb{H}_a^L$ is the L-fold Cartesian product of the complete allele graph over a alleles. In this case, the accessibility problem is of regular type, and it is fully specified by the scaled Hamming distance δ between the two endpoints. The exponential growth rate (16) of the expected number of quasi-accessible paths is then given by

$$\Gamma(\beta) = -\ln a - \beta + \delta \ln\left(e^{a\beta} - 1\right) + (1-\delta)\ln\left(e^{a\beta} + a - 1\right). \tag{20}$$

For the binary case $a = 2$ the threshold condition $\Gamma(\beta_c) = 0$ takes on the form

$$\sinh(\beta_c)^{\delta} \cosh(\beta_c)^{1-\delta} = 1, \tag{21}$$

which was first conjectured in [45], and subsequently proved in [47, 48]. For endpoints at maximal distance ($\delta = 1$), e^{β_c} is the solution of an algebraic equation of order a

$$\left(\mathrm{e}^{\beta_c}\right)^a - a\mathrm{e}^{\beta_c} - 1 = 0, \tag{22}$$

which can be solved in closed form for $a \leqslant 4$ (see table 1). For large a, the asymptotic expansions

$$\beta_c = \frac{\ln a}{a} + \frac{1+\ln a}{a^2} + \mathcal{O}\left(\frac{\ln a}{a^3}\right) \tag{23}$$

and

$$\beta_c \Gamma'(\beta_c) = \ln a + \frac{1+\ln a}{a} + \mathcal{O}\left(\frac{\ln a}{a^2}\right) \tag{24}$$

J. Stat. Mech. (2024) 034003

Table 1. Results for the complete allele graph with 2–4 alleles and endpoints at maximal distance $\delta = 1$. The last column shows the scaled asymptotic path length (19) at the threshold. In the biallelic case $a = 2$ the result for the path length was obtained independently in [49] for a related model. This work also contains additional results on the structure of accessible paths.

a	β_c	$\Gamma'(\beta_c)$	$\beta_c\Gamma'(\beta_c)$
2	$\arcsin(1) \approx 0.881$	$\sqrt{2} \approx 1.41$	≈ 1.25
3	$\ln\left(2\cos\frac{\pi}{9}\right) \approx 0.631$	$1 + 2\cos\left(\frac{2\pi}{9}\right) \approx 2.53$	≈ 1.82
4	$\ln\left(\frac{1}{\sqrt{2}}\right) + \sqrt{\sqrt{2} - \frac{1}{2}} \approx 0.509$	≈ 3.60	≈ 1.83

can be derived. Two general conclusions can be drawn from these expressions. First, HoC landscapes become much more accessible as the number of alleles increases, in the sense that $\beta_c \to 0$ for $a \to \infty$. Second, the increase of accessibility is achieved by paths that are only moderately longer (by a factor $\sim \ln a$) than the direct paths.

The results presented here and in [28] are in excellent agreement with earlier numerical simulations of accessible paths on Hamming graphs [42]. The authors of [42] focused on the probability of existence of at least one accessible maximal distance path ($d(\alpha,\omega) = L$) to the global fitness maximum. In this case, the fitness of the endpoint genotype is effectively $g(\omega) = 1$, whereas the starting point has uniformly distributed fitness. The probability of existence of such a path is then given by $1 - \beta_c$, and tends to unity for $a \to \infty$. For later reference we note that, because the fitness difference $\beta \in [-1,1]$ between two arbitrary genotypes α,ω has probability density $1 - |\beta|$, the probability for the existence of an accessible path $\alpha \to \omega$ between two randomly chosen genotypes at scaled distance δ is asymptotically given by

$$p_{\mathrm{rand}}(\delta) = \int_{-1}^{1} d\beta\,(1-|\beta|)\,\mathbb{P}[X_{\alpha,\omega} \geqslant 1|\beta,\delta] = \int_{\beta_c(\delta)}^{1} d\beta\,(1-\beta) = \frac{1}{2}\left[1-\beta_c(\delta)\right]^2. \tag{25}$$

4. Structured fitness landscapes

There is ample empirical evidence that real fitness landscapes are rugged, but they are not completely random [3, 6, 8]. Quantitative measures of ruggedness, such as the fraction of local peaks among genotype sequences, generally fall in between the limits of uncorrelated random and smooth (single-peaked) landscapes [22, 50]. In this section we examine the consequences that different kinds of landscape structure have on evolutionary accessibility. While it might be expected that reduced ruggedness implies increased accessibility, we will see that this is not generally true.

4.1. Geometry of the binary hypercube

Since theoretical work on models of structured fitness landscapes is largely restricted to the case of binary sequences, throughout this section we take the sequence space to

J. Stat. Mech. (2024) 034003

be the binary hypercube $\mathbb{H}_2^L$. To conveniently describe the geometry of this space, we introduce additional concepts that are specific to the binary case. First, the number

$$n_\sigma = \sum_{i=1}^{L} \sigma_i \tag{26}$$

of 1's in the sequence σ provides a natural one-dimensional coordinate that runs from $n_\sigma = 0$ for the 0-string $\sigma = (0,\ldots,0)$ to $n_\sigma = L$ for the 1-string $\sigma = (1,\ldots,1)$. Referring to the 0-string as the *wild type*, n_σ counts the number of mutations relative to the wild type, and the 1-string is the *full mutant*. A path in $\mathbb{H}_2^L$ is called *mutation-directed* (or *directed* for short) if n_σ changes monotonically along the path.

Second, it is often useful to represent a genotype σ by the subset of sites i in the *locus set* $\mathcal{L} = \{1,\ldots,L\}$ at which $\sigma_i = 1$ [51]. In this notation, the wild type is the empty set $\emptyset$, and the full mutant is identified with $\mathcal{L}$. An immediate corollary is that a mutation-directed path from α to ω exists iff α is a subset of ω or vice versa. Such a path is then also direct in the sense of section 2. Finally, the *antipode* $\bar{\sigma}$ of a binary sequence is obtained by reverting all the entries,

$$\bar{\sigma}_i = 1 - \sigma_i. \tag{27}$$

The antipode is the unique sequence at a maximal distance L from σ.

4.2. Additive landscapes

It is instructive to begin the discussion of structured landscapes with the seemingly trivial case of an additive landscape with a single peak [41]. The fitness function is of the form

$$g(\sigma) = \sum_{i=1}^{L} a_i \sigma_i, \tag{28}$$

and we assume for now that the coefficients a_i are positive. Then, the full mutant is the global and unique fitness peak, and all mutation directed paths to the peak are accessible. For a pair of genotypes α,ω, a directed path $\alpha \to \omega$ exists if

$$\alpha \subset \omega, \tag{29}$$

i.e. α is a subset of ω or (equivalently) ω is a *superset* of α. Thus, the peak is accessible from all genotypes ($\omega = \mathcal{L}$) and all genotypes can be accessed from the wild type ($\alpha = \emptyset$).

For two randomly chosen genotypes α,ω at Hamming distance d, the subset condition (29) requires that for all d sites at which the two sequences differ, the corresponding entries are given by $\alpha_i = 0$ and $\omega_i = 1$. This is true with probability $p_{\mathrm{rand}}(d) = 2^{-d}$, which (in contrast to the corresponding quantity (25) in the random case) tends to zero for $d \to \infty$. If (29) holds there are $d!$ different direct and directed paths from α to ω, and thus the expected number of accessible paths between random pairs of genotypes is given by

J. Stat. Mech. (2024) 034003

$$\mathbb{E}^{\text{add}}[X_{\alpha,\beta}|d] = 2^{-d}d! \tag{30}$$

which grows combinatorially for large d.

These statements remain valid if the coefficients a_i in (28) are random numbers without a definite sign, except that the unique and global peak is then given by the sequence with entries $\sigma_i = \frac{1}{2}[1+\text{sgn}(a_i)]$. The previous situation is recovered through a symmetry operation of the hypercube that maps this sequence to the full mutant. We conclude that for additive landscapes accessibility between arbitrary pairs of genotypes is low, but there is a direction of high accessibility, which generally depends on a_i. Of course, the low accessibility between random pairs is no obstacle for the evolutionary dynamics, since any strategy that accepts beneficial mutations and rejects deleterious ones will carry the population to the global fitness peak.

4.3. Kauffman's NK model

The NK model provides a simple framework for generating instances of fitness landscapes with a tunable degree of ruggedness [52, 53]. The basic idea is to write the fitness as a sum of contributions, each of which depends randomly on the variables at a subset of the locus set $\mathcal{L}$. The main tuning parameter of the model is the size of the subsets, which is generally assumed to be the same for all subsets and will be denoted here by k. When the subsets consist of single loci, $k=1$, the model reduces to the additive landscape (28) with random coefficients, whereas for $k=L$ the fitness is a random function of the entire sequence, and one recovers the HoC model discussed in section 3.

For a formal definition, we introduce the *interaction structure* [54] of the model as a collection of subsets $B_i \subset \mathcal{L}$ of size $|B_i| = k$, where $i = 1,\ldots,b$, and write the fitness function as

$$g(\sigma) = \sum_{i=1}^{b} \phi_i^{(k)}(\downarrow_{B_i}\sigma). \tag{31}$$

The map $\downarrow_{\mathcal{S}}: \mathbb{H}_2^L \to \mathbb{H}_2^{|\mathcal{S}|}$ projects the sequence σ onto the subspace spanned by the loci in the subset $\mathcal{S}$, and the $\phi_i^{(k)}$ are independent realizations of HoC fitness landscapes on $\mathbb{H}_2^k$. Examples of commonly used interaction structures are shown in figure 4.

The analysis of the NK model is straightforward for the case of the block structure (middle panel of figure 4) [55, 56]. We assume that L/k is an integer, and partition the locus set into $b = L/k$ disjoint interaction sets B_i, which will be referred to as blocks. Consider first the number of fitness peaks, which we denote by N. A sequence σ is a local fitness maximum iff each of the projected sequences $\downarrow_{B_i}\sigma, i = 1,\ldots,b$ is a maximum in its respective sub-landscape. Denoting the number of peaks of the ith sub-landscape by N_i, we thus have

$$N = \prod_{i=1}^{L/k} N_i. \tag{32}$$

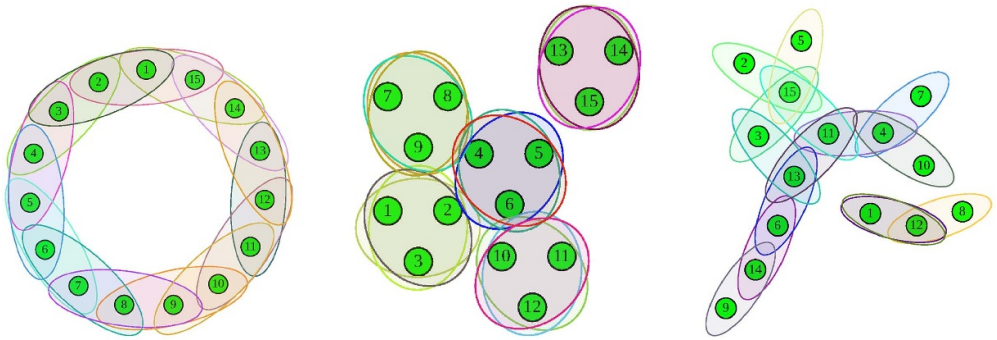

Figure 4. Interaction structures in the NK model. The figure illustrates possible choices of interacting subsets among $L=15$ loci. The loci enclosed in an ellipse belong to the same interaction set B_i. Note that the sets are not necessarily distinct. In the left and middle panel the subsets have size $k=3$, and in the right panel $k=2$. In the left panel the loci are arranged on a ring (*adjacent structure*), in the middle panel the subsets are either disjoint or identical (*block structure*), and in the right panel the subsets are chosen randomly (*random structure*). Adapted from [54], with permission from Springer Nature.

Since the sub-landscapes are independent, we can average over both sides to obtain

$$\mathbb{E}[N] = [\mathbb{E}(N_i)]^{L/k} = \left[\frac{2^k}{k+1}\right]^{L/k}, \tag{33}$$

where we have used that the mean number of peaks in an HoC landscape on $\mathbb{H}_a^L$ is [36]

$$\mathbb{E}^{\mathrm{HoC}}[N] = \frac{a^L}{(a-1)L+1}. \tag{34}$$

As expected, the expression (33) interpolates monotonically between the minimally rugged additive case ($k=1$) and the maximally rugged HoC landscape ($k=L$). The computation of the mean number of peaks for general interaction structures is non-trivial, but partial results are available and have been reviewed in [54]. Remarkably, the asymptotic behaviour of $\mathbb{E}[N]$ for large L is generally quite close to that predicted by the simple expression (33) for the block structure, and (33) has been conjectured to constitute an upper bound for general interaction structures [56].

A similar approach can be used to examine the evolutionary accessibility of the NK model with block interaction structure [56]. For simplicity, we focus on direct paths of length L that end at the global fitness maximum (henceforth denoted by ω) and start at its antipode $\alpha=\bar{\omega}$. The projection $\downarrow_{B_i}\omega$ of the global maximum to the ith block is the global maximum of the corresponding sub-landscape, and each step along the path has to increase the fitness contribution $\phi_i^{(k)}$ of the block in which the corresponding locus resides. Thus, each direct accessible path $\alpha\to\omega$ can be decomposed into accessible direct subpaths $\downarrow_{B_i}\alpha\to\downarrow_{B_i}\omega$. Whereas the ordering of the steps within a sub-landscape is fixed by the subpath, steps in different sub-landscapes can occur in arbitrary order.

J. Stat. Mech. (2024) 034003

Taken together, these considerations imply that the number of direct accessible paths can be written as

$$X_{\alpha,\omega} = \frac{L!}{(k!)^b} \prod_{i=1}^{b} X_{\downarrow_{B_i}\alpha,\downarrow_{B_i}\omega}, \tag{35}$$

where the combinatorial prefactor accounts for the different ways in which a given set of subpaths can be combined. Averaging both sides and making use of result (13) for the expected number of direct paths to the global maximum of an HoC landscape, we obtain

$$\mathbb{E}\left[X_{\alpha,\omega}\right] = \frac{L!}{(k!)^b}, \tag{36}$$

which again encompasses the limiting cases $k=1$ and $k=L$, and grows combinatorially fast with L for any fixed k.

On the other hand, for the global maximum of the full landscape to be at all accessible from the antipode, at least one accessible path has to exist in every sub-landscape. This implies that

$$\mathbb{P}\left[X_{\alpha,\omega} \geqslant 1\right] = \mathbb{P}\left[X_{\downarrow_{B_i}\alpha,\downarrow_{B_i}\omega} \geqslant 1\right]^{L/k}. \tag{37}$$

Since $\mathbb{P}[X_{\downarrow_{B_i}\alpha,\downarrow_{B_i}\omega} \geqslant 1] < 1$ for any $k > 1$, this quantity decays exponentially with L. Note that the exponential decay is much faster than the corresponding behaviour for direct paths in the random HoC model considered in section 3.1, where (when conditioning the endpoint to be of maximal fitness)

$$\mathbb{P}\left[X_{\alpha,\omega} \geqslant 1\right] = \mathbb{P}\left[\beta > \beta_c(L)\right] \approx \frac{\ln L}{L}. \tag{38}$$

These considerations can be generalized to indirect paths [54]. The length ℓ of a path covering the maximal distance L between the global fitness maximum and its antipode is a random variable, and similarly the number of (forward or backward) steps $\ell_i \geqslant k$ that the path spends in the ith sub-landscape is random, with $\sum_{i=1}^{b} \ell_i = \ell$. The multiplicity of each such path corresponding to the combinatorial prefactor in (35) is given by the multinomial coefficient

$$\binom{\ell}{\ell_1,\ldots,\ell_b} = \frac{\ell!}{\prod_{i=1}^{b} \ell_i!}. \tag{39}$$

Importantly, (37) remains valid and shows that the probability of the existence of an accessible path decays exponentially in L for any fixed k. The results described in section 3.3 show that, for large k, $\mathbb{P}[X_{\downarrow_{B_i}\alpha,\downarrow_{B_i}\omega}]$ converges to $1-\beta_c(\delta=1,a=2) \approx 0.119$. A nonzero limit of (37) is possible only if the number of interaction sets remains finite, which implies that k increases proportionally to L.

We conclude, therefore, that the accessibility of the global maximum in the NK landscape with block structure and fixed k is low, and remarkably similar to the accessibility

J. Stat. Mech. (2024) 034003

between random genotypes in the additive model discussed in section 4.2: The existence of paths is exponentially unlikely, but the expected number of paths grows combinatorially fast with the path length. Refined arguments presented in [54, 57] show that the low accessibility of the block model is in fact typical of most common NK interaction structures, including the adjacent and random structures displayed in figure 4. The precise statement is that, for interaction structure graphs that satisfy a certain local boundedness condition, accessible paths with $d(\alpha,\omega) = L$ are exponentially unlikely for $L \to \infty$. This behaviour is strikingly different from that of uncorrelated random landscapes, where such paths exist with certainty provided the fitness difference between the endpoints is sufficiently large.

4.4. Rough Mount Fuji landscapes

The example of the NK model shows, somewhat counterintuitively, that the reduced ruggedness of certain types of structured landscapes can decrease their accessibility. An example of a landscape that displays the expected negative correlation between ruggedness and accessibility is the Rough Mount Fuji (RMF) model, defined as a weighted linear superposition of an additive and a random (HoC) landscape [58, 59],

$$g(\sigma) = \xi_\sigma + cd\left(\sigma,\sigma^*\right). \tag{40}$$

Here, ξ_σ are continuous i.i.d. random variables, $c > 0$ a constant and σ^* a reference sequence. For $c = 0$ the landscape is random, and when c is sufficiently large (larger than the fluctuations of ξ_σ) it becomes additive with a unique global peak at the antipode $\overline{\sigma^*}$ of the reference sequence. While the number of peaks decreases monotonically with increasing c [59], the probability of existence of direct accessible paths from σ^* to $\overline{\sigma^*}$ tends to unity for $L \to \infty$ for any $c > 0$ [43].

5. Highly rugged yet highly accessible fitness landscapes

Empirical evidence suggests that real fitness landscapes may be both highly rugged and highly accessible. For example, a recent study of a fitness landscape comprising almost all $4^9 = 262\,144$ combinations of nucleotide mutations at nine sites in the *folA* gene in *E. coli* identified 514 peaks among the 18 019 sequences that gave rise to functional proteins [26]. To put this number into perspective, we note that, according to (34), the probability for a randomly chosen genotype in a HoC fitness landscape with a alleles and sequence length L to be a fitness peak is $\frac{1}{(a-1)L+1}$; this is simply the probability for the fitness of the peak genotype to be the largest among $(a-1)L+1$ continuous i.i.d. random variables [36]. Since most of the functional sequences of the *folA* landscape have the maximal number of $(a-1)L = 27$ functional neighbours, a completely random landscape would have on average $18\,019/28 \approx 644$ peaks, only slightly more than the experimental value. Thus, by any measure, the empirical landscape is highly rugged. However, at the same time, the 73 highest fitness peaks (those whose fitness value exceeds the wild type fitness) were highly accessible, in the sense that their basins of attraction comprised a significant fraction of all genotypes, with a median value of 69%.

J. Stat. Mech. (2024) 034003

Here and in the following, the *basin of attraction* (BoA) of a fitness peak consists of all genotypes from which the peak can be reached via an accessible path. Note that this definition differs from the *gradient basins* considered in previous work, where each point in a landscape is uniquely assigned to the peak that is reached by a greedy (steepest ascent) walk starting from the point [21, 53, 60, 61]. Whereas gradient basins partition the genotype space into disjoint sets, the BoAs considered here can (and do) overlap. Like accessible paths, they depend only on the orientation of the fitness graph induced by the landscape, and not on the actual fitness values.

In this section we describe a recently discovered class of fitness landscape models that combine a high degree of ruggedness with high accessibility. These landscapes are characterized by a certain accessibility property (AP) that we define in the next subsection. The AP was first derived for a model of evolution in changing environments motivated by the evolution of antibiotic resistance [51, 62], and this context is explained in section 5.2. In section 5.3 we provide a simple sufficient condition for the AP, which allows us to devise a minimal representative for this class of models and to explore some of its properties (section 5.4). Throughout, we consider landscapes defined on the binary hypercube $\mathbb{H}_2^L$. The generalization of these concepts to multiallelic fitness landscapes with $a > 2$, which is required for applying them to the empirical landscape of [26] discussed above, remains a task for future research (see section 6).

5.1. The accessibility property

For the formulation of the property of interest we make use of the set notation introduced in section 4.1. Importantly, this implies that the landscapes we consider are *graded*, in the sense that their properties vary systematically along the coordinate n_σ connecting the wild type to the full mutant. In this respect they differ from the HoC and NK-models discussed previously, which are statistically isotropic (averaged over the random variables defining the landscape). By contrast, the RMF landscapes considered in section 4.4 are similarly structured along the coordinate defined by the distance to the reference sequence.

We say that a fitness landscape has the *subset–superset AP* if the following holds [51]:

Any peak genotype is accessible from all its subsets and supersets via all direct paths.

This implies that the basins of attraction of fitness peaks have a simple structure that is reminiscent of the additive landscapes of section 4.2. The BoA of a peak genotype σ with n_σ mutations consists of two subcubes spanned by σ and $\emptyset$ (the subsets of σ) and by σ and $\mathcal{L}$ (the supersets of σ), respectively. All $n_\sigma!$ (direct and directed) paths $\emptyset \to \sigma$ are accessible, as well as all $(L-n_\sigma)!$ paths $\mathcal{L} \to \sigma$, and the size S_σ of the BoA is bounded from below by

$$S_\sigma \geqslant 2^{n_\sigma} + 2^{L-n_\sigma} - 2. \tag{41}$$

Here we adopt the convention that a genotype does not belong to its own BoA. Note that by definition, $S_\sigma \geqslant L$. An example of a fitness graph with two peaks that satisfies the AP is shown in figure 5.

J. Stat. Mech. (2024) 034003

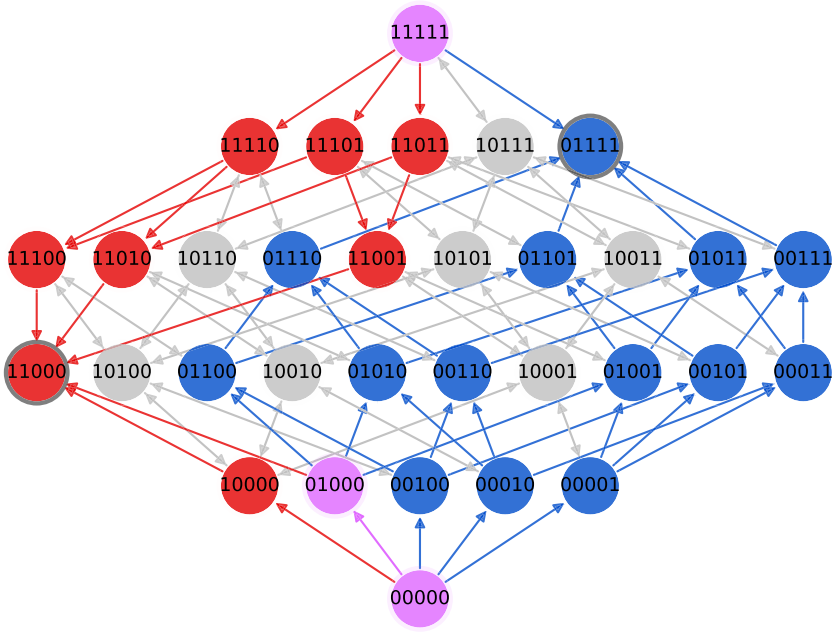

Figure 5. Illustration of the subset–superset AP. The fitness graph of dimension $L = 5$ has two peak genotypes, (0,1,1,1,1) and (1,1,0,0,0), which are highlighted by grey circles. The direction of the coloured arrows is forced by the AP, whereas the orientation of the light-grey arrows is arbitrary. Genotypes that are sub- or supersets of (0,1,1,1,1), and thus must belong to the corresponding BoA, are coloured blue, and the genotypes that must belong to the BoA of (1,1,0,0,0) are red. Purple genotypes are sub- or supersets of both peaks.

Next, we derive an upper bound on the number of peaks N that a landscape satisfying the AP (an AP landscape for short) can have [63]. Since no peak can be in the basin of attraction of another peak, different peak genotypes cannot be subsets or supersets of each other. According to Sperner's theorem, the largest collection of subsets of a set of S elements that has this property is of size $\binom{S}{\lfloor S/2 \rfloor}$ [64], and we conclude that

$$N \leqslant \binom{L}{\lfloor L/2 \rfloor}. \tag{42}$$

The bound is saturated by a fitness landscape where all genotypes with $n_\sigma = \lfloor L/2 \rfloor$ are peaks and no other peaks exist. For large L the bound (42) is of order $\frac{2^L}{\sqrt{\pi L/2}}$, which differs from the maximal number of peaks in an unconstrained landscape, 2^{L-1} [65, 66], only by a polynomial factor.

5.2. Tradeoff-induced fitness landscapes

The model of tradeoff-induced fitness landscapes (TIL model) describes the evolutionary processes of a population facing a tradeoff in its adaptation to two extreme environments [51, 62]. The model generates a family of fitness landscapes parametrized

J. Stat. Mech. (2024) 034003

by a single, continuous environmental parameter. The scenario guiding the construction of the model is the evolution of antibiotic resistance in bacteria, where the environmental parameter determining the strength and direction of selection is the drug concentration. In this context, the *dose–response curve* $g(\sigma, c)$ represents the growth rate of a bacterial population with genotype σ exposed to the drug concentration c [67]. Motivated by the empirical findings laid out in [51], the dose–response curves are assumed to be of the form

$$g(\sigma, c) = r_\sigma f(c/m_\sigma). \tag{43}$$

The function $f(x)$ is monotonically decreasing, accounting for the suppression of bacterial growth with the increasing drug concentration. Setting $f(0) = 1$, the parameter r_σ is identified as the *drug-free growth rate* of the strain, whereas m_σ defines the concentration scale at which the growth rate is significantly suppressed. In other words, m_σ quantifies the *resistance* of the strain. With increasing x, the function $f(x)$ may either reach zero growth rate at some finite value x_c, in which case $c = m_\sigma x_c$ is called the minimal inhibitory concentration, or $f(x) \to 0$ asymptotically for large x. A commonly used functional form of the latter type is the Hill function

$$f(x) = \frac{1}{1 + x^h} \tag{44}$$

with a Hill coefficient $h > 0$.

It follows from (43) and (44) that strains with large r_σ are favoured at low concentrations, whereas strains with large m_σ have an advantage at high concentrations. A tradeoff between these two extremes occurs if strains with large r_σ have small m_σ and vice versa. A simple (and empirically supported) way of implementing such a tradeoff is to write the drug-free growth rate and the resistance in the form

$$r_\sigma = \prod_{i=1}^{L} r_i^{\sigma_i}, \quad m_\sigma = \prod_{i=1}^{L} m_i^{\sigma_i}, \tag{45}$$

where $r_i < 1$ and $m_i > 1$ are parameters associated with the ith mutation that can be measured experimentally or drawn from some suitable probability distribution [51]. Equation (45) implies that genotypes with many mutations (large n_σ) have, on average, a low drug-free growth rate and high resistance.

As a consequence of the tradeoff between the drug-free growth rate and resistance, the dose–response curves of different strains may cross [51]. If the crossing occurs between strains whose genotypes are mutational neighbours, the corresponding arrow in the fitness graph flips. Under suitable (mild) conditions on the scaling function $f(x)$ and the parameters r_i and m_i, it can be ensured that the dose–response curves of neighbouring genotypes cross exactly once [51]. In this way, the TIL model defined by equations (43)–(45) induces an evolution on the set of fitness graphs, which is parametrized by c.

J. Stat. Mech. (2024) 034003

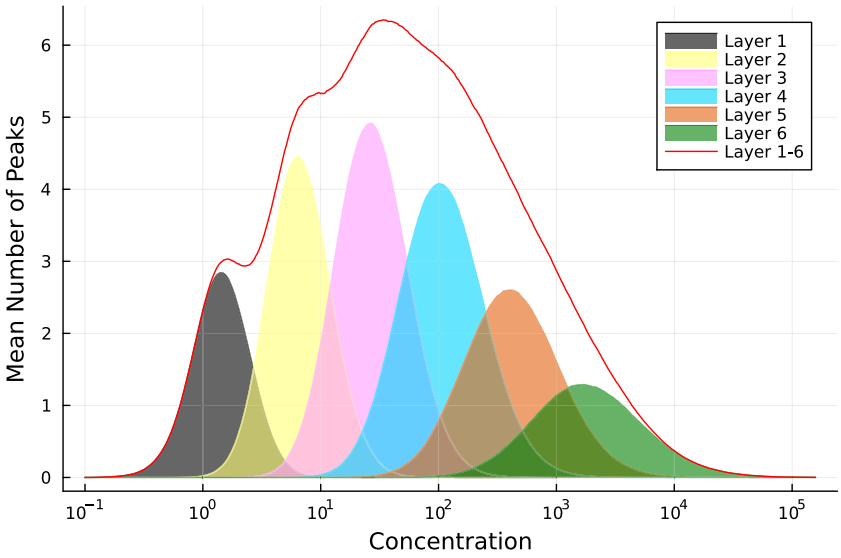

Figure 6. Fitness maxima in tradeoff-induced landscapes. The figure shows the mean number of fitness maxima in the TIL model as a function of concentration c, obtained from 15 000 realizations with $L = 7$. The parameters $\{r_i\}_{i=1,\dots,L}$ and $\{m_i\}_{i=1,\dots,L}$ were sampled from the joint distribution described in [51] and the concentration was varied from 10^{-1} to 10^{8} in 1000 logarithmically spaced steps. The shaded curves show the number of peaks in a particular *layer* of sequence space, where the mutation number n_σ is fixed, and the red line shows the total number of peaks in layers 1–6. At very small and very large concentrations, the fitness landscape is single-peaked with the unique fitness maximum in layer $n_\sigma = 0$ and $n_\sigma = 7$, respectively. These peaks are not included in the figure.

For small and large c the graphs are simple, because the corresponding fitness landscapes are multiplicative:

$$g(\sigma, 0) = r_\sigma, \quad g(\sigma, c \to \infty) \approx \frac{r_\sigma m_\sigma^h}{c^h} = c^{-h} \prod_{i=1}^{L} \left(r_i m_i^h\right)^{\sigma_i}. \tag{46}$$

In these limits the fitness landscape is single-peaked, and all direct paths to the peak are accessible[3]. However, at intermediate concentrations, the tradeoff induces a significant amount of ruggedness, and it can be shown that the number of peaks grows exponentially with L [51, 62]. The genotypes that optimize the tradeoff (and therefore have the highest fitness) at concentration c have on the order of

$$n_{\text{opt}} \approx \frac{\ln c}{\langle \ln m_i \rangle} \tag{47}$$

mutations, and the ruggedness is maximal when $n_{\text{opt}} \approx \frac{L}{2}$ (figure 6). The AP explained in section 5.1 arises in the TIL model as a consequence of the ordering of the concentration

[3] If g is multiplicative, $\ln g$ is additive in the sense of section 4.2.

J. Stat. Mech. (2024) 034003

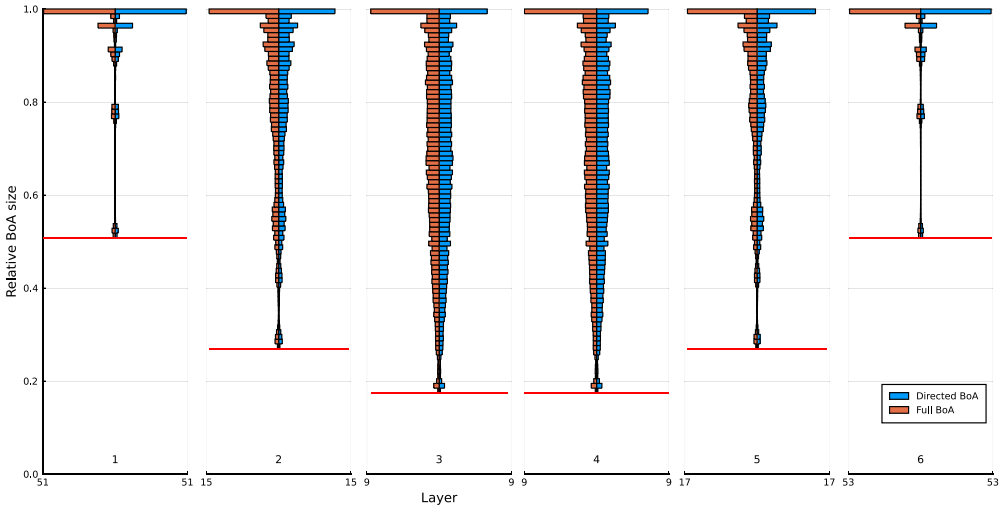

Figure 7. Basins of attraction of fitness maxima in tradeoff-induced landscapes. The figure shows histograms for the relative basin size (48) obtained from simulations of the TIL model with $L=7$ (see figure 6 for details). The peak genotypes are grouped along the x-axis into layers of equal mutation number. Only instances are shown where neither the wild type ($n_\sigma=0$) nor the full mutant ($n_\sigma=7$) are peaks, because in these cases the landscape must be single-peaked. For each peak genotype, the *maximal* size of its BoA attained at any concentration is recorded. In the vertical histograms, left orange bars refer to the relative size of the full BoAs spanned by all accessible paths, and right blue bars to the relative size of the BoA's spanned by directed paths. Note the different horizontal scale of each vertical histogram pair. The red horizontal lines mark the lower bound (41) divided by $2^L-2=126$. This normalization accounts for the fact that, if a peak has a BoA of minimal size, there must be at least one other peak that accommodates the remaining genotypes. As a consequence, the maximal number of non-peak genotypes is 2^L-2.

values at which the dose–response curves of neighbouring genotypes cross, and it holds at any concentration (see [51, 62] for details).

In [51], the accessibility of the TIL landscape was quantified through the reachability of the global fitness peak in adaptive walk simulations, and it was found to be much higher than for an NK landscape of comparable ruggedness. Here, we explore the accessibility of the landscape by measuring the sizes of the BoAs of fitness peaks (figure 7). The relative basin size S_σ^{rel} depicted in the figure is obtained by normalizing the basin size S_σ by the number of non-peak genotypes,

$$S_\sigma^{\mathrm{rel}}=\frac{S_\sigma}{2^L-N}. \tag{48}$$

The figure shows that BoAs that are considerably larger than the lower bound (41) are common, and in fact many have maximal size. Inspection of the data shows that most of these instances are not single-peaked landscapes (for which, trivially, $S_\sigma^{\mathrm{rel}}=1$).

J. Stat. Mech. (2024) 034003

Interestingly, the histograms representing the full and the directed BoAs have very similar shapes.

5.3. Universal epistasis

'Epistasis', like 'invertebrate', is a term that really means 'everything else'
Sean H Rice [68]

In modern evolutionary genetics, the term *epistasis* is used to describe any kind of interaction in the effects of different mutations on a phenotype or on fitness [3, 8, 12, 69–73]. In the present context, this means that the fitness function $g(\sigma)$ deviates from the additive form (28). As an illustrative example, we consider two mutations occurring at positions i and j in a sequence σ. For notational convenience we introduce the mutation operator $\Delta_l : \mathbb{H}_2^L \to \mathbb{H}_2^L$, which flips the allele at locus l and is formally defined through [54]

$$(\Delta_l \sigma)_k = (1-\delta_{kl})\,\sigma_k + \delta_{kl}\,(1-\sigma_k)\,. \tag{49}$$

The effect sizes of the two single mutations are given by $s_{i,j} = g(\Delta_{i,j}\sigma) - g(\sigma)$, and the pairwise epistatic interaction ϵ_{ij} measures the deviation of the effect size of the double mutant from the additive prediction $s_i + s_j$:

$$\epsilon_{ij} = g\,(\Delta_i\,\Delta_j \sigma) - g\,(\sigma) - (s_i + s_j) = g\,(\Delta_i\,\Delta_j \sigma) + g\,(\sigma) - g\,(\Delta_i\,\sigma) - g\,(\Delta_j \sigma)\,. \tag{50}$$

We say that the mutations at positions i and j interact epistatically in the *background* σ if $\epsilon_{ij} \neq 0$, and the interaction is positive or negative according to the sign of ϵ_{ij}. Similar expressions can be written to quantify epistatic interactions of higher order [38, 50, 74].

The term *universal epistasis* was introduced in [66] to describe a global constraint on a fitness landscape that relates the ordering between the fitness effects of a mutation in different backgrounds to the inclusion relation (in the sense of section 4.1) between the background genotypes. Specifically, we consider two genotypes σ, σ' with $\sigma' \subset \sigma$, as well as a set τ of sites not contained in σ. Then the fitness landscape displays universal positive (negative) epistasis, if the fitness effect of the composite mutational event $\sigma \to \sigma \cup \tau$ is always greater (smaller) than the effect of the mutation $\sigma' \to \sigma' \cup \tau$. In the present context, the condition of universal negative epistasis (UNE) is of primary interest, which can be formally written as

$$g\,(\sigma \cup \tau) - g\,(\sigma) \leqslant g\,(\sigma' \cup \tau) - g\,(\sigma') \quad \text{for any } \sigma' \subset \sigma,\ \tau \subset \mathcal{L} \setminus \sigma. \tag{51}$$

In the appendix we show that the condition (51) holds for arbitrary σ, σ', τ provided it is satisfied for all epistatic interactions between pairs of single mutations in all backgrounds. In this special case, $d(\sigma, \sigma') = 1$ and the set τ in (51) consists of a single element.

The importance of the relation (51) lies in the observation (first noted in [51]) that a fitness landscape with purely negative epistasis satisfies the AP defined in section 5.1. Here, we show how the AP follows from (51). Suppose that the genotype σ is a fitness peak, and consider a subset genotype $\sigma' \subset \sigma$. We then need to show that the addition

J. Stat. Mech. (2024) 034003

of any mutation $i \in \sigma \setminus \sigma'$ to σ' increases fitness. Because σ is a peak, we know that $g(\sigma) - g(\sigma \setminus \{i\}) > 0$ for any i, and using (51) it follows that

$$g(\sigma' \cup \{i\}) - g(\sigma') \geqslant g(\sigma) - g(\sigma \setminus \{i\}) > 0. \tag{52}$$

The accessibility of the peak from superset genotypes can be proved in the same way.

5.4. Fitness landscapes with an intermediate phenotype

A simple way to generate fitness landscapes with the UNE property (51) is to combine a linear genotype–phenotype map with a suitably chosen nonlinear function that maps the phenotype to fitness [7]. Fitness landscapes based on such a construction have been studied theoretically in the context of Fisher's geometric model [75–78], but are also used in the analysis and interpretation of empirical data sets [24, 79–81]. Conceptually, they constitute an attempt to partially reinstate the intermediate phenotype in the genotype–fitness map (figure 1).

For the case of a one-dimensional intermediate phenotype z, the fitness function is of the form

$$g(\sigma) = \Phi[z(\sigma)] \quad \text{with} \quad z(\sigma) = \sum_{i=1}^{L} a_i \sigma_i. \tag{53}$$

To proceed, we make two further assumptions.

(i) The function Φ is concave, that is,

$$\Phi(x+a) - \Phi(x) < \Phi(y+a) - \Phi(y) \quad \text{for any } a > 0,\ x > y. \tag{54}$$

(ii) The coefficients a_i are positive.

For genotypes σ, σ', τ with $\sigma' \subset \sigma$ and $\tau \subset \mathcal{L} \setminus \sigma$ it follows from (ii) that $z(\sigma) > z(\sigma')$ and $z(\sigma \cup \tau) = z(\sigma) + z(\tau)$, $z(\sigma' \cup \tau) = z(\sigma') + z(\tau)$. Using additionally the concavity property (i), we find that

$$\begin{aligned} g(\sigma \cup \tau) - g(\sigma) &= \Phi[z(\sigma) + z(\tau)] - \Phi[z(\sigma)] \\ &< \Phi[z(\sigma') + z(\tau)] - \Phi[z(\sigma')] = g(\sigma' \cup \tau) - g(\sigma') \end{aligned} \tag{55}$$

and (51) is established. As an example, a landscape that saturates the bound (42) on the number of peaks is obtained by taking $a_i \equiv \bar{a} > 0$ to be a constant and choosing Φ as a non-monotonic function with its maximum at $x = \bar{a}\lfloor \frac{L}{2} \rfloor$. Figure 8 illustrates the construction for the general case of unequal phenotypic effects a_i, and figures 9 and 10 show results of an exploratory numerical study of the model using exponentially distributed coefficients a_i. Comparing figures 7 and 10, we see that the size distributions of the BoAs of fitness peaks have much less weight at values close to $S_\sigma^{\mathrm{rel}} \approx 1$ in the present model. Importantly, the AP, which is satisfied for both models, does not fully specify the properties of the BoAs.

J. Stat. Mech. (2024) 034003

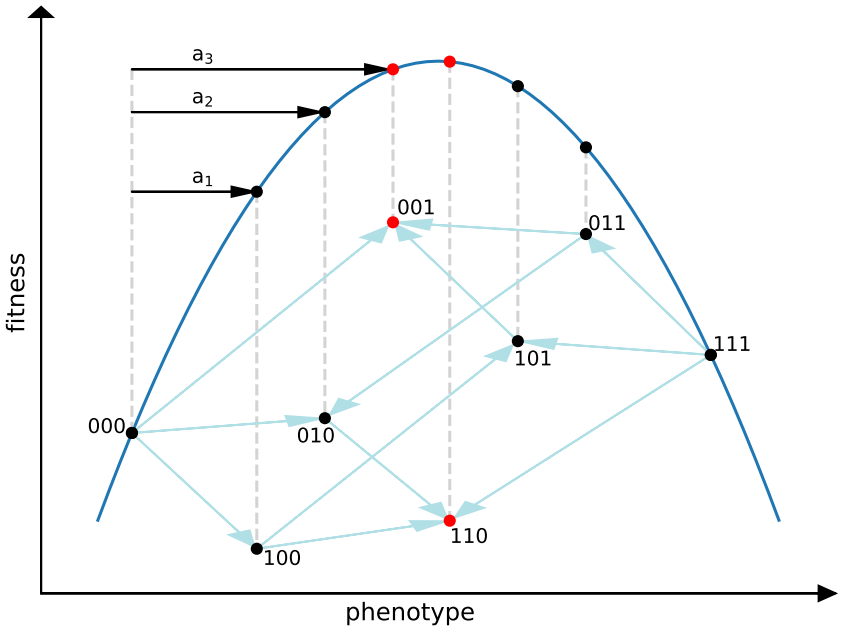

Figure 8. Genotypic fitness landscape induced by a one-dimensional nonlinear phenotype–fitness map. The figure illustrates the construction of a fitness landscape of the form (53) with $L=3$. The fitness of a genotype depends nonlinearly on the corresponding phenotype value, which is a linear combination of the positive mutational effects a_1, a_2, a_3. The dashed horizontal lines connect the fitness values to the nodes of the 3-cube. Because the phenotype–fitness map is non-monotonic, the genotypic fitness landscape can have multiple peaks, which are marked in red.

In fact, the positivity condition (ii) on the a_i can be relaxed. For general (nonzero) coefficients we define a transformation $\sigma \to \tilde{\sigma}$ through

$$\tilde{\sigma}_i = \begin{cases} \sigma_i & a_i > 0 \\ 1-\sigma_i & a_i < 0 \end{cases} \tag{56}$$

and rewrite the linear phenotype as

$$z(\sigma) = \sum_{i=1}^{L} a_i \sigma_i = \sum_{i:a_i>0} |a_i| \sigma_i - \sum_{i:a_i<0} |a_i| \sigma_i = \sum_i |a_i| \tilde{\sigma}_i + \sum_{i:a_i<0} a_i. \tag{57}$$

Thus z is a linear function of the $\tilde{\sigma}$ with positive coefficients, and the relation (55) holds in the transformed coordinates.

These considerations can be applied to Fisher's geometric model with a one-dimensional phenotype [77, 78]. In this model, the coefficients a_i are random variables sampled from a symmetric distribution with zero mean, and the phenotype–fitness map has a unique maximum at $z=0$. This corresponds to a situation where selection favours an intermediate value of a phenotype, a scenario that makes good sense biologically [82].

J. Stat. Mech. (2024) 034003

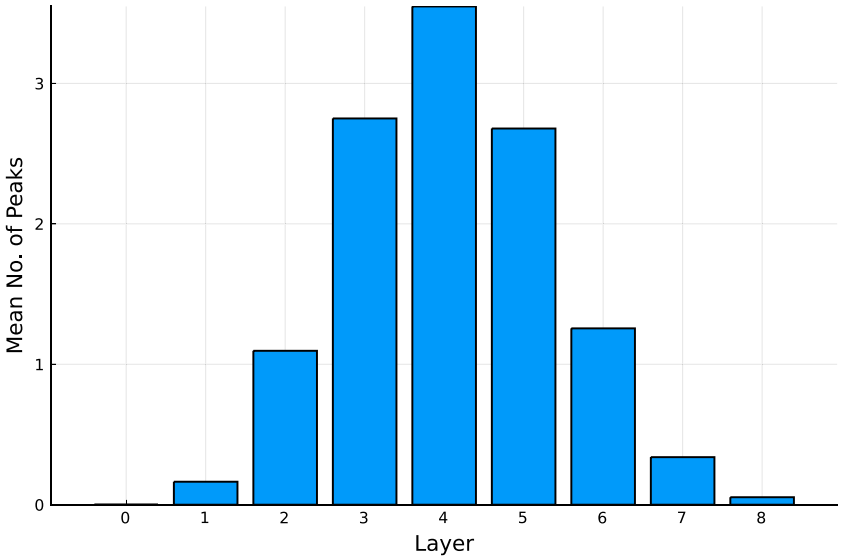

Figure 9. Distribution of the number of fitness peaks in a landscape with universal negative epistasis. The figure shows the distribution of the number of fitness peaks across layers of constant mutation number n_σ for a fitness landscape of the form (53). The dimension is $L=8$, the coefficients $a_i > 0$ are chosen from an exponential distribution with unit mean, and the phenotype–fitness map is $\Phi(z) = -(z - L\langle a_i \rangle/2)^2 = -(z-4)^2$. The data constitute an average over 20 000 realizations.

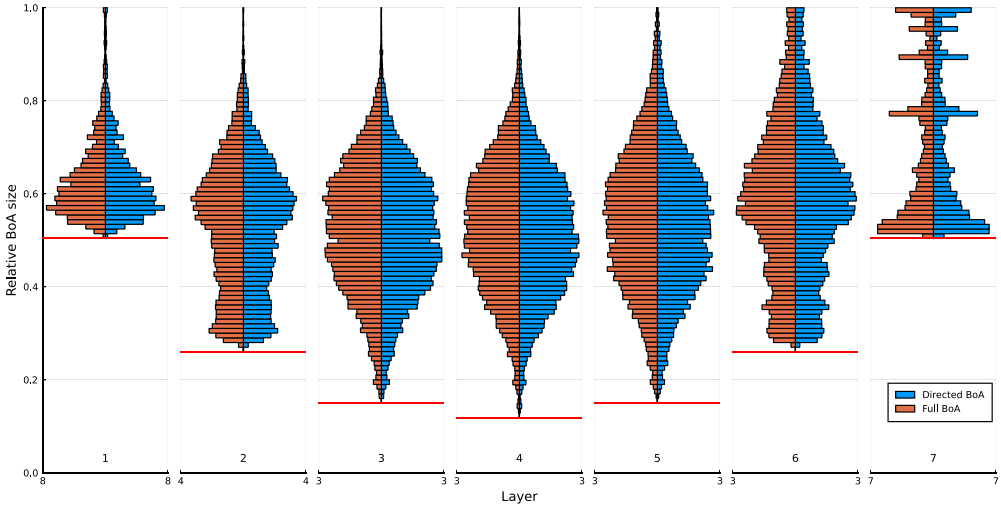

Figure 10. Basins of attraction of fitness maxima in a landscape with universal negative epistasis. The figure shows histograms for the relative size (48) of the BoAs of fitness peaks for the model described in figure 9. The layout is the same as in figure 7.

J. Stat. Mech. (2024) 034003

In [78] it was shown that the expected number of fitness peaks in the one-dimensional model is of the order of

$$\mathbb{E}[N] \sim \frac{2^L}{L^{3/2}} \tag{58}$$

for large L, and the full distribution of N was computed. Thus, these fitness landscapes are highly rugged. The arguments presented above show that, provided the function $\Phi(z)$ is concave, the model satisfies the subset–superset AP. A simple example of a concave function is the inverse parabola, $\Phi(z) \sim -z^2$. As explained in [78], with this choice, Fisher's geometric model is closely related to the antiferromagnetic Hopfield model [83], and it can also be mapped to the number partitioning problem [84]. This suggests that the AP holds for the optimization landscapes of these problems as well.

6. Summary and outlook

The dynamics of many complex systems can be described as a search process in a high-dimensional landscape defined by some objective function such as energy or fitness. Probabilistic models defined over discrete configuration spaces provide a versatile framework for exploring the properties of these landscapes, and are of great current interest in diverse fields ranging from evolutionary biology to deep learning [85]. In this article we have reviewed a particular notion of landscape accessibility that originated in the context of biological fitness landscapes. Detailed case studies of two classical fitness landscape models, the HoC and the NK model, were summarized in sections 3 and 4. Somewhat surprisingly, we find that accessible pathways covering distances linear in the landscape dimension L are likely to exist even in the uncorrelated random HoC landscape, provided the fitness gain covered by the path is large enough. At the same time, correlated landscapes of the NK type are generally less, not more accessible than random landscapes.

Section 5 describes a novel class of fitness landscapes with a graded structure that enables high ruggedness (many peaks) to coexist with high accessibility. These landscapes are characterized by the subset–superset AP, which couples the accessibility of fitness peaks to the inclusion relations between genotypes when interpreted as sets of loci. An immediate but striking consequence of the AP is a *lower* bound on the size of the basin of attraction of any fitness peak that grows exponentially with L, as $2^{L/2}$ (see equation (41)). The AP was originally found in a model of tradeoff-induced fitness landscapes [51], but it can also be derived as a consequence of negative epistatic interactions. Importantly, the minimal (necessary) conditions for the AP to hold are not presently known.

We close by describing some problems for future research that need to be addressed in order to apply the ideas presented here to large-scale empirical data sets such as those in [26, 27]. First, the work on structured fitness landscapes should be generalized to sequence spaces with more than two alleles [86]. In particular, the considerations for the

J. Stat. Mech. (2024) 034003

AP landscapes discussed in section 5 need to be modified to account for the geometry of multiallelic sequence spaces. Second, a better understanding of the correlations between different quantifiers of landscape topography for single instances of a probabilistic model is required. For example, how do the number and height of fitness peaks correlate with the sizes of their basins of attraction? While properties of BoAs remain unexplored even for the most basic fitness landscape models, recent studies on the joint distribution of the number of peaks and the frequency of different types of pairwise epistatic interactions [63, 87, 88] have begun to shed light on related questions.

Acknowledgments

We are much indebted to Kristina Crona, Suman Das, Jasper Franke, Muhittin Mungan, Stefan Nowak and Benjamin Schmiegelt for their contributions to the results described here, and to Suman Das and Muhittin Mungan for comments on the manuscript. Many thanks to Peter Mörters for making us aware of the connection to submodularity. This work was supported by DFG within the Grants SFB 680, SFB 1310 and SPP 1590.

Note added

After submission of the revised version of the manuscript, we noticed that universal negative epistasis is equivalent to a property of set functions known as submodularity [89]. The consequences of this observation will be investigated elsewhere.

Appendix. Pairwise negative epistasis implies universal negative epistasis

For the following considerations it is useful to rewrite the UNE condition (51) in the equivalent form [66]

$$g(\upsilon\cup\upsilon')+g(\upsilon\cap\upsilon')\leqslant g(\upsilon)+g(\upsilon')\,, \tag{A.1}$$

which is obtained from (51) by setting $\upsilon=\sigma$ and $\upsilon'=\sigma'\cup\tau$. In this formulation, the condition holds for arbitrary pairs υ,υ'. We start from two genotypes at maximal distance, $d(\upsilon,\upsilon')=L$, and show that (A.1) is satisfied, provided it holds for pairs at strictly smaller distance. By iterating the argument, it follows that it is sufficient to demand non-positive epistasis for all pairs of genotypes at distance $d=2$. Geometrically, these pairs make up the 2-faces of the binary L-cube.

The maximal distance condition $d(\upsilon,\upsilon')=L$ implies that $\upsilon'=\bar{\upsilon}$ and therefore $\upsilon\cup\upsilon'=\mathcal{L}$ and $\upsilon\cap\upsilon'=\emptyset$. We now choose a proper, non-empty subset $\upsilon_1\subset\upsilon'$, which can be done in $2^{|\upsilon'|}-2$ different ways. If this is not possible because $|\upsilon'|<2$, υ_1 is chosen to be a subset of υ instead, and the roles of υ' and υ are interchanged. Next define $\upsilon_2\equiv\upsilon_1\cup\upsilon$. Our strategy is to assume UNE for the pairs

$$(\upsilon_1,\upsilon_1')\equiv(\upsilon_1,\upsilon)\quad\text{and}\quad(\upsilon_2,\upsilon_2')\equiv(\upsilon_2,\upsilon')=(\upsilon_1\cup\upsilon,\upsilon')\,. \tag{A.2}$$

The following identities are easily checked:

$$\upsilon_1\cap\upsilon_1'=\emptyset,\ \ \upsilon_1\cup\upsilon_1'=\upsilon_1\cup\upsilon,\ \ \upsilon_2\cap\upsilon_2'=\upsilon_1,\ \ \upsilon_2\cup\upsilon_2'=\mathcal{L}. \tag{A.3}$$

Moreover, making use of the general relation

$$d(\chi,\chi') = d(\chi' \cup \chi, \chi' \cap \chi) \tag{A.4}$$

for arbitrary pairs of genotypes χ,χ', we find

$$d(\upsilon_1,\upsilon_1') = d(\emptyset,\upsilon_1\cup\upsilon) = |\upsilon_1\cup\upsilon| < L,\;\; d(\upsilon_2,\upsilon_2') = d(\mathcal{L},\upsilon_1) = L - |\upsilon_1| < L \tag{A.5}$$

because υ_1 is a proper, non-empty subset of $\upsilon' = \bar{\upsilon}$. Thus the distances between the genotypes making up the pairs (A.2) are strictly smaller than L. The UNE conditions for the pairs read

$$g(\upsilon_1\cup\upsilon_1') + g(\upsilon_1\cap\upsilon_1') \leqslant g(\upsilon_1) + g(\upsilon_1') \tag{A.6}$$
$$g(\upsilon_2\cup\upsilon_2') + g(\upsilon_2\cap\upsilon_2') \leqslant g(\upsilon_2) + g(\upsilon_2')\,. \tag{A.7}$$

Adding the two inequalities and cancelling identical terms on both sides, the UNE condition (A.1) for (υ,υ') follows. Unless the pairs (A.2) already form 2-faces, $d(\upsilon_1,\upsilon_1') = d(\upsilon_2,\upsilon_2') = 2$, the argument is repeated for the smaller subcubes spanned by $(\upsilon_1\cup\upsilon_1', \upsilon_1\cap\upsilon_1')$ and $(\upsilon_2\cup\upsilon_2', \upsilon_2\cap\upsilon_2')$. The same proof applies to the case of universal positive epistasis considered in [66].

References

[1] Wright S 1932 The roles of mutation, inbreeding, crossbreeding and selection in evolution *Proc. 6th Int. Congress of Genetics* vol 1 pp 356–66
[2] Svensson E I and Calsbeek R 2012 *The Adaptive Landscape in Evolutionary Biology* (Oxford University Press)
[3] de Visser J A G M and Krug J 2014 *Nat. Rev. Genet.* **15** 480
[4] Hartl D L 2014 *Curr. Opin. Microbiol.* **21** 51–57
[5] Kondrashov D A and Kondrashov F A 2015 *Trends Genet.* **31** 24–33
[6] Fragata I, Blanckaert A, Louro M A D, Liberles D A and Bank C 2019 *Trends Ecol. Evol.* **34** 69–82
[7] Manrubia S *et al* 2021 *Phys. Life Rev.* **38** 55–106
[8] Bank C 2022 *Annu. Rev. Ecol. Evol. Syst.* **53** 457–79
[9] Provine W B 1986 *Sewall Wright and Evolutionary Biology* (The University of Chicago Press)
[10] Agarwala A and Fisher D S 2019 *Theor. Pop. Biol.* **130** 13–49
[11] Greenbury S F, Louis A A and Ahnert S E 2022 *Nat. Ecol. Evol.* **6** 1742–52
[12] Weinreich D M, Watson R A and Chao L 2005 *Evolution* **59** 1165–74
[13] Poelwijk F J, Kiviet D J, Weinreich D M and Tans S J 2007 *Nature* **445** 383
[14] Carneiro M and Hartl D L 2010 *Proc. Natl. Acad. Sci. USA* **107** 1747–51
[15] Franke J, Klözer A, de Visser J A G M and Krug J 2011 *PLoS Comput. Biol.* **7** e1002134
[16] Weinreich D M, Delaney N F, DePristo M A and Hartl D L 2006 *Science* **312** 111–4
[17] DePristo M A, Hartl D L and Weinreich D M 2007 *Mol. Biol. Evol.* **24** 1608–10
[18] Lozovsky E R, Chookajorn T, Brown K M, Imwong M, Shaw P J, Kamchonwongpaisan S, Neafsey D E, Weinreich D M and Hartl D L 2009 *Proc. Natl. Acad. Sci. USA* **106** 12025–30
[19] Palmer A C, Toprak E, Baym M, Kim S, Veres A, Bershtein S and Kishony R 2015 *Nat. Commun.* **6** 7385
[20] Bank C, Matuszewski S, Hietpas R T and Jensen J D 2016 *Proc. Natl. Acad. Sci. USA* **113** 14085–90
[21] Wu N C, Dai L, Olson C A, Lloyd-Smith J O and Sun R 2016 *eLife* **5** e16965
[22] Aguilar-Rodríguez J, Payne J L and Wagner A 2017 *Nat. Ecol. Evol.* **1** 0045
[23] Domingo J, Diss G and Lehner B 2018 *Nature* **558** 117–21
[24] Pokusaeva V O *et al* 2019 *PLoS Genet.* **15** e1008079
[25] Moulana A, Dupic T, Phillips A M, Chang J, Nieves S, Roffler A A, Greaney A J, Starr T N, Bloom J D and Desai M M 2022 *Nat. Commun.* **13** 7011
[26] Papkou A, Garcia-Pastor L, Escudero J A and Wagner A 2023 *Science* **382** eadh3860

J. Stat. Mech. (2024) 034003

[27] Westmann C A, Goldbach L and Wagner A 2023 *bioRxiv Preprint* (https://doi.org/10.1101/2023.08.25.554764) (posted online 27 August 2023, accessed 31 August 2023)
[28] Schmiegelt B and Krug J 2023 *J. Math. Biol.* **86** 46
[29] Stadler P F and Happel R 1999 *J. Math. Biol.* **38** 435–78
[30] Altenberg L 2015 (arXiv:1508.07866)
[31] de Visser J A G M, Park S-C and Krug J 2009 *Am. Nat.* **174** S15–S30
[32] Crona K, Greene D and Barlow M 2013 *J. Theor. Biol.* **318** 1–10
[33] Plotkin J B and Kudla G 2011 *Nat. Rev. Genet.* **12** 32–42
[34] Zwart M P, Schenk M F, Hwang S, Koopmanschap B, de Lange N, van de Pol L, Nga T T T, Szendro I G, Krug J and de Visser J A G M 2018 *Heredity* **121** 406–21
[35] Kingman J F 1978 *J. Appl. Probab.* **15** 1–12
[36] Kauffman S and Levin S 1987 *J. Theor. Biol.* **128** 11–45
[37] Derrida B 1981 *Phys. Rev.* B **24** 2613–2326
[38] Crona K, Gavryushkin A, Greene D and Beerenwinkel N 2017 *eLife* **6** e28629
[39] Nowak S and Krug J 2013 *Europhys. Lett.* **101** 66004
[40] Krug J 2021 Accessibility percolation in random fitness landscapes *Probabilistic Structures in Evolution* (EMS Press) pp 1–22
[41] Schmiegelt B 2023 Structure and accessibility of fitness landscapes *PhD Thesis* University of Cologne (available at: http://kups.ub.uni-koeln.de/id/eprint/64747)
[42] Zagorski M, Burda Z and Waclaw B 2016 *PLoS Comput. Biol.* **12** e1005218
[43] Hegarty P and Martinsson A 2014 *Ann. Appl. Probab.* **24** 1375–95
[44] Berestycki J, Brunet E and Shi Z 2016 *Bernoulli* **22** 653–80
[45] Berestycki J, Brunet E and Shi Z 2017 *ALEA Lat. Am. J. Probab. Math. Stat.* **14** 45–62
[46] Martinsson A 2018 *Ann. Probab.* **46** 1004–41
[47] Martinsson A 2015 (arXiv:1501.02206)
[48] Li L 2018 *J. Theor. Probab.* **31** 2072–111
[49] Kistler N and Schertzer A 2020 (arXiv:2012.04076)
[50] Szendro I G, Schenk M F, Franke J, Krug J and de Visser J A G M 2013 *J. Stat. Mech.* P01005
[51] Das S G, Direito S O, Waclaw B, Allen R J and Krug J 2020 *eLife* **9** e55155
[52] Kauffman S A and Weinberger E D 1989 *J. Theor. Biol.* **141** 211–45
[53] Weinberger E D 1991 *Phys. Rev.* A **44** 6399–413
[54] Hwang S, Schmiegelt B, Ferretti L and Krug J 2018 *J. Stat. Phys.* **172** 226–78
[55] Perelson A S and Macken C A 1995 *Proc. Natl. Acad. Sci.* **92** 9657–61
[56] Schmiegelt B and Krug J 2014 *J. Stat. Phys.* **154** 334–55
[57] Schmiegelt B 2016 Sign epistasis networks *Master's Thesis* University of Cologne (available at: http://kups.ub.uni-koeln.de/id/eprint/10398)
[58] Aita T, Uchiyama H, Inaoka T, Nakajima M, Kokubo T and Husimi Y 2000 *Biopolymers* **54** 64–79
[59] Neidhart J, Szendro I G and Krug J 2014 *Genetics* **198** 699–721
[60] Wolfinger M T, Svrcek-Seiler W A, Flamm C, Hofacker I L and Stadler P F 2004 *J. Phys. A: Math. Gen.* **37** 4731–41
[61] Franke J and Krug J 2012 *J. Stat. Phys.* **148** 705–22
[62] Das S, Krug J and Mungan M 2022 *Phys. Rev.* X **12** 031040
[63] Oros D 2022 Hypercubes, peak patterns and universal positive epistasis *Master's Thesis* University of Cologne (available at: http://kups.ub.uni-koeln.de/id/eprint/63386)
[64] Sperner E 1928 *Math. Z.* **27** 544–8
[65] Haldane J 1931 *Math. Proc. Camb. Phil. Soc.* **27** 137–42
[66] Crona K, Krug J and Srivastava M 2023 *J. Math. Biol.* **86** 62
[67] Regoes R R, Wiuff C, Zappala R M, Garner K N, Baquero F and Levin B R 2004 *Antimicrob. Agents Chemother.* **48** 3670–6
[68] Rice S H 2000 The evolution of developmental interactions: epistasis, canalization and integration *Epistasis and the Evolutionary Process* (Oxford University Press) pp 82–98
[69] Wolf J B, Brodie E D III and Wade M J (eds) 2000 *Epistasis and the Evolutionary Process* (Oxford University Press)
[70] Phillips P C 2008 *Nat. Rev. Genet.* **9** 855–67
[71] Ferretti L, Schmiegelt B, Weinreich D, Yamauchi A, Kobayashi Y, Tajima F and Achaz G 2016 *J. Theor. Biol.* **396** 132–43
[72] Domingo J, Baeza-Centurion P and Lehner B 2019 *Annu. Rev. Genom. Hum. Genet.* **20** 17.1–17.28

J. Stat. Mech. (2024) 034003

[73] Krug J 2021 Epistasis and evolution *Oxford Bibliographies in Evolutionary Biology* (Oxford University Press)
[74] Poelwijk F J, Krishna V and Ranganathan R 2016 *PLoS Comput. Biol.* **12** e1004771
[75] Fisher R A 1958 *The Genetical Theory of Natural Selection* (Dover)
[76] Tenaillon O 2014 *Annu. Rev. Ecol. Evol. Syst.* **45** 179–201
[77] Hwang S, Park S-C and Krug J 2017 *Genetics* **206** 1049–79
[78] Park S-C, Hwang S and Krug J 2020 *J. Phys. A: Math. Theor.* **53** 385601
[79] Rokyta D, Joyce P, Caudle S B, Miller C, Beisel C J and Wichman H A 2011 *PLoS Genet.* **7** e1002075
[80] Schenk M F, Szendro I G, Salverda M L M, Krug J and de Visser J A G M 2013 *Mol. Biol. Evol.* **30** 1779–87
[81] Otwinowski J, McCandlish D M and Plotkin J B 2018 *Proc. Natl. Acad. Sci. USA* **115** E7550–8
[82] Srivastava M and Payne J L 2022 *PLoS Comput. Biol.* **18** e1010524
[83] Nokura K 1998 *J. Phys. A: Math. Theor.* **31** 7447–59
[84] Ferreira F and Fontanari J F 1998 *J. Phys. A: Math. Theor.* **31** 3417–28
[85] Mézard M 2023 *Indian J. Phys.* (https://doi.org/10.1007/s12648-023-03029-8)
[86] Srivastava M, Rozhoňová H and Payne J L 2023 *J. Phys. A: Math. Theor.* **56** 455601
[87] Saona R, Kondrashov F A and Khudiakova K A 2022 *Bull. Math. Biol.* **84** 74
[88] Riehl M, Phillips R, Pudwell L and Chenette N 2022 *J. Phys. A: Math. Theor.* **55** 434002
[89] Krause A and Golovin D 2014 Submodular function maximization *Tractability: Practical Approaches to Hard Problems* (Cambridge University Press) pp 71–104

J. Stat. Mech. (2024) 034003